# DDP – A tool for life-cycle risk management

**Steven L. Cornford**
Strategic Systems Technology Program Office
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
(818)354-1701
steven.cornford@jpl.nasa.gov

**Martin S. Feather**
Quality Assurance Office
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
(818)354-1194
martin.feather@jpl.nasa.gov

**Kenneth A. Hicks**
Strategic Systems Technology Program Office
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
(818)354-0595
kenneth.hicks@jpl.nasa.gov

*Abstract* - At JPL we have developed, and implemented, a process for achieving life-cycle risk management[*]. This process has been embodied in a software tool and is called Defect Detection and Prevention (DDP). The DDP process can be succinctly stated as: determine where we want to be, what could get in the way and how we will get there. The 'determine where we want to be' is captured as trees of requirements and the 'what could get in the way' is captured as trees of potential failure modes[†]. Scoring the impacts of these failure modes on the requirements results in a prioritized set of failure modes. The user then selects from a set of PACTs (Preventative measures, Analyses, process Controls and Tests) each of which has an effectiveness versus the various failure modes. It is the goal of the DDP process to optimally select the subset of the PACTs[‡] which minimizes the residual risk subject to the project resource constraints.

The DDP process is intended to facilitate risk management over the entire project life cycle beginning with architectural and advanced technology decisions all the way through operation. As the project design, technology content and implementation approach matures, the requirements and failure mode trees are elaborated upon to accommodate the additional information. Thus, the DDP process is a systematic, continuous, top-down approach to managing risk. Implementation of the DDP process requires a critical mass of expertise (usually the project team and a few specialists) and captures both their engineering judgement as well as available quantitative data. This additional data may result from models, layouts, prototype testing, other focused risk evaluations and institutional experiences. The DDP process also identifies areas where additional information would be advantageous, thus allowing a project to target critical areas of risk or risk uncertainty. This also allows the project to identify those areas which would benefit the most from application of other quantitative tools and methods (e.g. Monte Carlo simulations, FMECAs, fault trees).

The software tool supports the DDP process by providing guidance for implementing the process steps, graphical visualizations of the various trees, their inter-relationships and the current risk landscape. The tool is capable of supporting on-the-fly knowledge elicitation as well as integrating off-line deliberations. There are a variety of available outputs including graphs, trees and reports as well as clear identification of the driving requirements, 'tall-pole' residual risks and the PACTs which have been selected and agreed upon. The DDP process has been applied at various levels of assembly


[†] Failure mode here is used in its most general sense – inability to achieve the requirements.
[‡] Note that each PACT has some resource costs associated with if (e.g. dollars, schedule, mass).

including the system and subsystem levels, as well as down to the component level. Recently significant benefits have been realized from application to advanced technologies, where the focus has been on increasing the infusion rates of these technologies by identification and mitigation of risks prior to delivery to a project.

## TABLE OF CONTENTS



## 1. INTRODUCTION

We have developed and implemented a process [1,2] embodied in a software tool [3,4,5]. The DDP process can be very briefly summarized as determining: Where we want to be, what could get in the way, how we will get there. This is a very general statement and is the primary reason that the DDP process has been, or can be, applied from the mission suite architecture level all the way down to the interconnections on a bare die. It has been said that DDP could be applied to develop approaches to handling teen-age children, but this is probably beyond the scope of any known process.

*Background*

Within NASA and the aerospace industry as a whole, there has long been a general feeling among project personnel that too many PACTs are being performed and that the 'value-added' of many of them is limited. However, it is unclear exactly where the excess is located. Within the Mission Assurance community, there is a general worry that either not enough or barely enough is being done. However, it is unclear where the 'barely screened' failure modes are located. Both of these views are probably correct as the next figure illustrates.

*"Screening out the defects"*

As a background for understanding the DDP underlying process, consider **Figure 1**. This picture is intended to illustrate that there are many potential mission failure modes (these are depicted as solid lines). Various PACTs are implemented (depicted as boxes with text in them) and these PACTs detect or prevent the occurrence of some of the failure modes (depicted as the disappearance of a solid line). However, these PACTs may have limited scope (the width of the box does not cover all possible solid lines) and non-perfect effectiveness ('escapes' or un-detected or un-prevented failure modes are depicted as dashed lines). Examining the chart further, we see some cases of interest. The solid line at the far right was completely unscreened and has impacted the mission success. The solid line in the middle had one chance of being screened (System Test) and if that PACT is not performed or is modified in some way that failure mode would not caught. There are also numerous cases where the boxes, or portions of them, were redundant, and did not all need to be performed. In the language of the picture, it is the goal of the DDP process is to facilitate the selection of the optimal combination of boxes which are consistent with the project requirements.

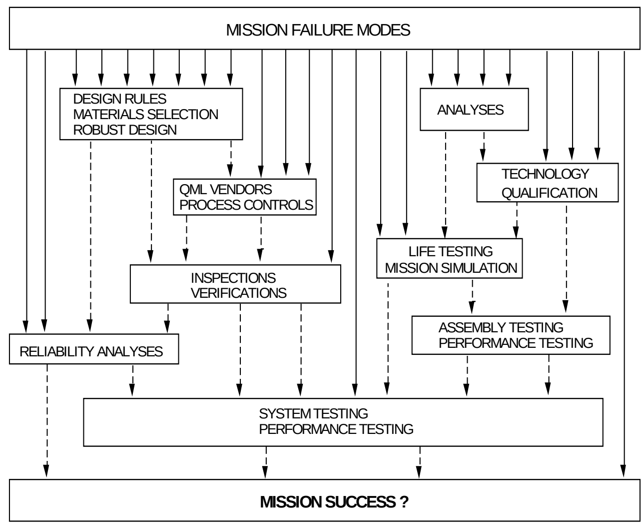


**Figure 1** 'Waterfall' chart illustrating how potential failure modes (solid lines) are detected or prevented) by the application of the PACTs in the boxes and 'escapes' are shown as dashed lines. Note that this figure is not to scale.

*The DDP Process*

As noted above, the DDP process is intended to help the projects choose the optimal set of PACTs consistent with their resource constraints. We do not want to select PACTs based on 'what we've always done' but rather on their relative effectiveness at detecting or preventing the failure modes which are relevant to the specific project application. This requires knowledge of what the potential failure modes are and their relative importance. It also requires knowledge about what potential PACTs are available and their relative effectiveness and resource costs.

The DDP process can be summarized as consisting of 4 steps shown in **Figure 2**. These four steps will now be described.

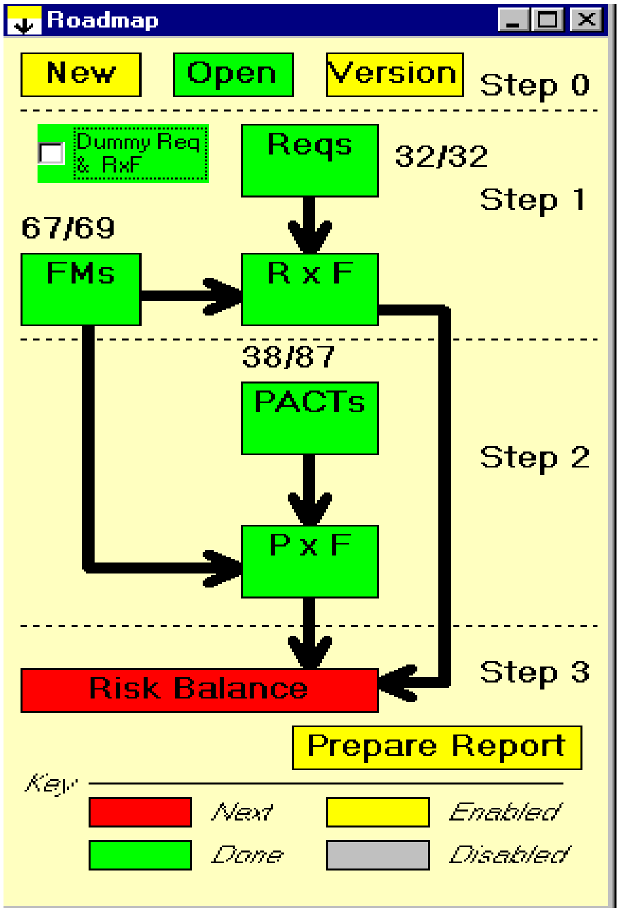


**Figure 2** Roadmap from the DDP tool illustrating the 4 principal steps in implementing the DDP process. In this particular evaluation, 32 Requirements were utilized (out of 32 identified), 67 of 69 Failure Modes were deemed relevant and 38 of 87 possible PACTs were selected for implementation.

*Step 0 (Understand the technology)*– The goal of this first step is to develop a more detailed understanding of the product/technology under evaluation. Prior to any DDP evaluation, it is imperative that the DDP team understands the product under evaluation, be it an entire spacecraft or a particular technology[§]. Thus, the first step is an information exchange meeting where the DDP team is brought up to speed via any and all available information (e.g. documents, drawings, block diagrams, layouts, test results). This also allows the team to be fine-tuned by adding an expert or two to achieve a 'critical mass' of expertise. For example, spacecraft-level evaluations would require more systems and architectural expertise, while detailed technology evaluations would require more materials and device operation expertise.

The outputs of the DDP process will only be as thorough as the available information and the breadth of the involved experts. The level of evaluation fidelity will only be as good the level of information detail. At higher levels of evaluation, the scoring will be primarily based on engineering judgement and corporate knowledge bases. At lower levels of evaluation, more detailed information is generally available (testing and modeling results, etc.) and scoring can use 'more digits of accuracy'. The DDP process is designed to use all available information and intentionally allows the mixing of engineering judgement and quantitative analytical results.

*Step 1 (Develop the Requirements matrix)*- The goal of this next step is to develop a prioritized set of failure modes (or risk elements). In the DDP process this is accomplished by completing a Requirements Matrix (**R**), shown in Error! Reference source not found., which determines the relative importance of the failure modes and the extent to which the various requirements are driving the risk. The impact of each failure mode on each requirements is scored[**] as the percentage of the requirement lost should the failure mode occur. Summing down the columns (weighted by the relative importance of each requirement and the a priori likelihood of each failure mode) yields the criticality of each failure mode. Summing across the rows yields the extent to which each requirement is at risk and thus identifies driving requirements which may then be reexamined for their necessity.

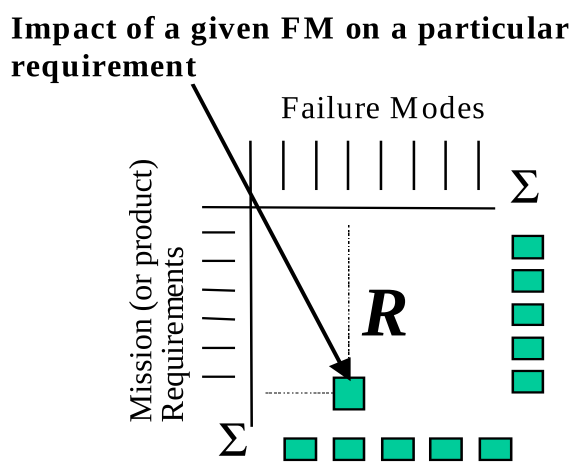


**Figure 3** The Requirements matrix maps the impacts of each failure mode (should it occur) on each requirement. Note that this schematic does not show the tree structures of both the requirements and the failure modes.

It is important to note that within the DDP process, risk criticality can be modified by changing the requirements or their relative importance. The extent to which the

---

[§] This step is shorter or longer depending on the extent to which the DDP team consists of members of the 'already-up-to-speed' project team.

[**] At higher levels of evaluation, the information regarding effectiveness and impact is much more intuitive and scoring usually utilizes numbers like 0, 0.1, 0.3, 0.9 and 1.0. As the level of insight increases, the scoring can utilize more precision as the evaluation moves into the more analytical realm.

requirements are driving the design is valuable information early in the requirement definition phase of the project life cycle, where the project may be developing requirements early in the life cycle based on the intuition from scientists or project engineers.

*Step 2 (Develop the Effectiveness matrix)*– The goal of this step is to develop a set of options (PACTs) for preventing or detecting the failure modes. First the relative effectiveness of the various PACT options versus the relevant failure modes. Some PACTs may be very effective against a few specific failure modes (e.g. PC board delamination testing versus PC board delamination issues), while others may be less effective against a much broader spectrum of failure modes (e.g. end-to-end functional testing versus all possible functional failure modes).

The Effectiveness matrix ( E ), shown in **Figure 5**, captures the effectiveness of each PACT against each failure mode and is scored as the fractional reduction in the likelihood of occurrence of the failure mode. In addition, each PACT has resource costs associated with it. At the end of this step, the DDP prepares the team with a collection of possible PACTs (sometimes a baseline has already been chosen by the project prior to the evaluation) and the costs associated with them.

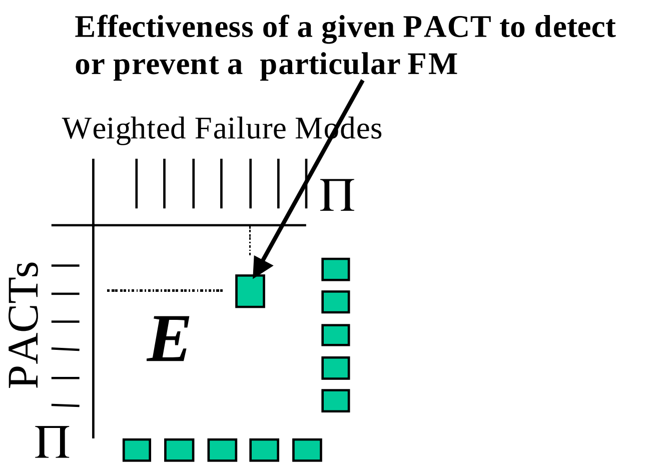


**Figure 5** The Effectiveness matrix maps the probability of detecting or preventing each failure mode by each PACT should the PACT be implemented. Note that this schematic does not show the tree structures of both the PACTs and the failure modes.

*Step 3 (Balance the risk and iterate)*– The goal of this last step, which is the primary goal of the DDP process, is to select the optimal subset of PACTs which achieves balanced risk consistent with the project resource constraints. The risks have now all been appropriately weighted so a bar chart of risks shows the tallest bars are those which are the most important to address. The DDP process is graphically summarized in **Figure 5**.

By choosing (or not) various PACT combinations, one may watch the result ripple through the risk balance and the resource totals. Thus, PACTs being applied to already low enough risk elements may be un-selected to free up the resources necessary to apply other PACTs to those risk elements which are currently too high. While the goal of the process is to balance risk consistent with project programmatic and resource constraints, it is possible that there is no set of PACTs which accomplishes this goal. For example, extensive use of new technology, complex operational scenarios, and overbearing cost and schedule constraints may not allow a minimal risk program to be achieved. However, in this case the DDP process will allow the project to identify specific areas where less aggressive requirements, additional resources or additional accepted risk would make a difference in achieving the goals. The process of selecting the optimal subset of PACTs is facilitated by a variety of tool features, but the current version of the tool does not have an automatic optimization algorithm, although this is currently under development.

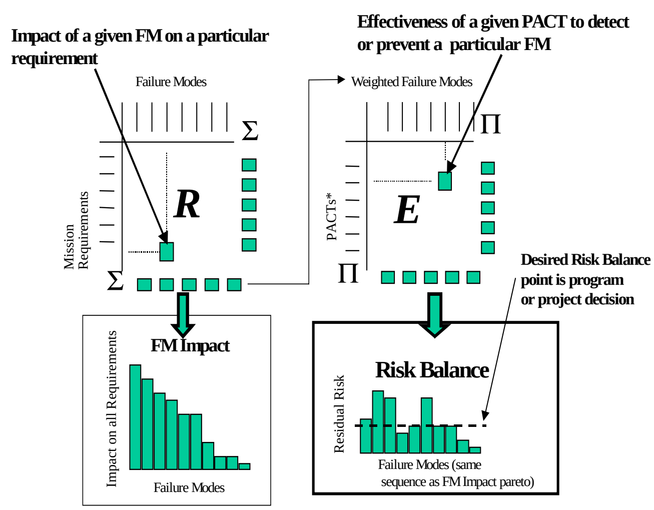


**Figure 5** Graphical summary of the DDP process. The Requirements matrix results in a prioritized set of Failure modes (sorted in the bar chart), a representation of the driving requirements (not shown), and the residual risk if the selected PACTs are implemented.

The DDP tool allows project personnel to zoom in and out of any of these bar charts in order allow risk examination at a variety of levels. For example, a project manager may only want to know which subsystem is the most at risk, while the subsystem engineer may wish to know portion of the subsystem is at most risk, etc.

## USING THE DDP PROCESS TO DO RISK MANAGEMENT

*DDP as part of the NASA Risk Management Process*

Risk management has some well-defined phases and objectives which NASA has carefully articulated [6]. The main elements can be summarized as:

- Risk Identification
- Risk Analysis
- Risk Planning
- Risk Tracking
- Risk Control

Using all of these elements has been demonstrated to be essential to the successful management of risk and are accommodated within the DDP process.

- Risk Identification

In order to manage risk, one needs to know what the risks are. They may result from a variety of sources including the technology content, environmental interactions, the implementation and operation approaches, programmatic constraints and the mission duration. The DDP process begins with articulating the requirements ("where we are trying to go") and then utilizes available project information, experts and brainstorming to develop an initial tree of potential failure modes. This tree is pruned and shaped by the results of the evaluation and other introduced information (e.g. Fault Trees, FMECAs). This tree development and shaping process integrates top-down, system-level risk (failure mode) identification with bottoms-up risk (failure mechanism) identification in an attempt to more systematically develop a achieve a more complete failure mode tree. This process of tree evolution allows the DDP process to stay in phase with the evolving design and implementation decisions of the project.

- Risk Analysis

The DDP process analyzes the consequences of the potential risks (failure modes) by scoring their impact on the requirements should they occur. This results in a requirement-driven risk list where failure modes are derive their criticality from their impact on (possibly) weighted requirements. Note that the failure modes may also be weighted by a likelihood of occurrence should nothing be done (lightning strikes are an obvious example of a potential failure mode which may not occur even if nothing is done). This risk analysis also results in prioritized list of driving requirements so it is immediately visible which requirements are at risk and to what extent.

- Risk Planning

The number of possible PACTs available for implementation far exceeds the resources of any project. Furthermore, different PACTs have a different effectiveness against different failure modes. In most cases, there are also a number of PACTs available for each failure mode (e.g. design rules, process controls, testing, modeling, inheritance). The DDP tool allows the users to identify combination of PACTs which will not only adequately address the risk but also most effectively utilize the available project resources (e.g. mass may be at a premium but additional schedule may be available). Users can explore the possibilities of implementing different APCT suites, adopt baselines and examine a variety of 'what-if' options.

- Risk Tracking

The tool has a number of report formats which can be used by different personnel for different reasons. The risks can be listed and include the PACTs which were selected to ameliorate their impacts. These PACTs may be then examined in detail to ensure the adequacy of the overall mitigation approach. In addition, the user can generate a report of the PACT selected and what they were intended to prevent or detect. This allows the PACT engineers to design the PACT to focus on the specific reasons for it's performance and to avoid implementing portions of them for which no real benefit is expected. There are a number of other reports available.

- Risk Control

If requirements (or their relative importance) change, or PACTs which were planned for implementation are not performed, or new potential failure modes are discovered, the tool performs real-time modification of the resultant risk so the project always has an up-to-date top N risk list. This allows the project team to effectively control risk and watch it's growth or decline as the design evolves and the results of implementation become available.

*Using the DDP process over the entire project life cycle*

DDP is intended to be used as part of a continuous risk management process. The requirements, failure modes, articles and PACTs are represented as tree structures to emphasize the ability to use an evolving hierarchal approach to perform the DDP assessments. This allows the project team risk management efforts to remain as current as possible with the rapidly evolving project design. The early assessments will be more high-level, on average, as the team is assessing the risk from high-level information such as architectures, block and state diagrams, etc. The result of this high-level evaluation is to target general areas of risk with general approaches to mitigation. For example, a particularly complex architecture may result in increased emphasis on model development and interface management while the reliance on heritage hardware may result in an increased emphasis on ensuring consistent manufacturing and evaluating the similarity of the intended usage environments.

At higher levels of evaluation, the scoring is more intuitive and becomes more analytical at lower-level, more detailed evaluations. For very detailed evaluations of specific issues (probability of occurrence, impact on requirements, effectiveness of PACTs), a wide variety of additional analytical tools are more suited to generating the greater levels of accuracy. In this way, the DDP process can help guide the projects in where to focus their high-powered analytical capability. Thus, DDP attempts to integrate the best of the intuitive and analytical approaches in the regimes where each has the highest fidelity (see **Figure 6**).

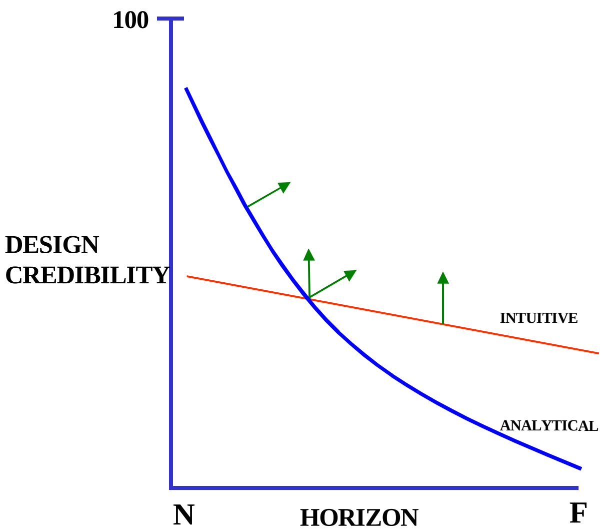


**Figure 6** DDP integrates intuitive and analytical approaches where each has the highest level of fidelity. When looking at more distant horizons, the information is more vague and intuitive insight is generally more credible, while for near horizons where a plethora of detailed data is available, analytical approaches produce the highest credibility. The arrows represent directions for improvement in these two approaches.

The DDP process described can be applied anywhere in the project life cycle – it just uses different information (and team members) depending on the level of the evaluation. This facilitates risk management over the entire life cycle beginning with architectural design and technology content, continuing through fabrication and integration, and finishing all the way out at the end of operation. **Figure 7** illustrates this range of applicability along with other products which may be relevant.

*Value and Applicability of the DDP process*

DDP has been shown to be a systematic, top-down approach to manage risk, which also incorporates bottoms up information as evaluations by a team focused on the details may uncover failure modes and mechanisms which can have system level implications.

This ability to incorporate the evolving design into the tool is a key part of its value and applicability. As an example, at a higher level of evaluation, an architectural approach may have the greatest potential for reducing risk and when this preventative measure (P in PACT) is adopted by the project it results in a collection of derived requirements which may be at risk from more specific failure modes, which might be mitigated by a collection of more specific PACTs, and so on.

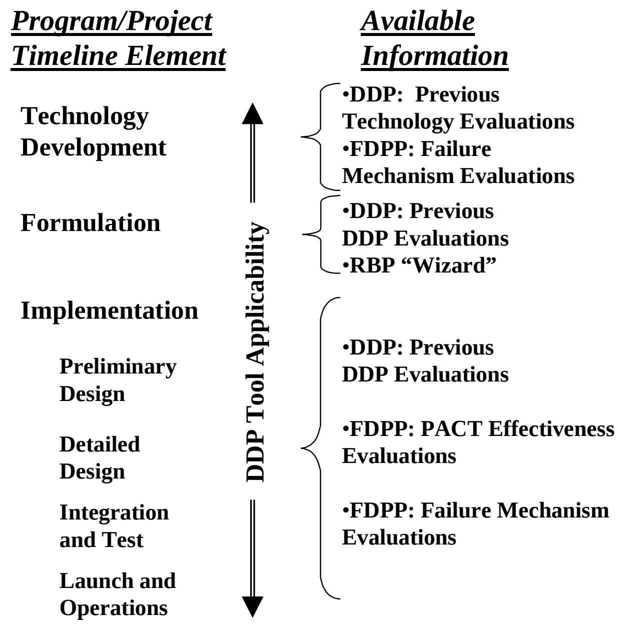


**Figure 7** DDP applicability as a function of project life cycle. Different types of information may more useful in different phases of application. FDPP=NASA Code Q Failure Detection and Prevention Program and RBP=Risk Balancing Profiles.

Another benefit of the DDP process is that it works best when the DDP team consists of project team members and required additional specialists. This helps ensure that the project maintains ownership of the answer and avoids problems with bringing external teams up to speed.

While the process can start with a clean sheet of paper (no previously identified requirements, failure modes or PACTs), we have seen that the *effectiveness of PACTs can be evaluated on* collections of failure mechanisms without specific knowledge of which project will be utilizing the data. NASA Code Q is funding both the tool development and tool population efforts and a variety of products regarding PACT effectiveness and the failure modes associated with various technology types are available[7]. This default information may be imported directly into an evaluation and modified as required.

The process also helps the project identify areas where additional information would be valuable, either to address a risk element or to reduce the uncertainty of a risk element. Thus, the project can identify areas which benefit the most from other tools applied.

This allows the project to always maintain a current target list of critical risk or risk uncertainties. Note that the DDP process weights risk elements by impact on weighted requirements (which are a direct flow-down from the mission success criteria and their relative importance!) to ensure that the project team members are focusing on those issues which most affect the success of the mission.

## THE DDP SOFTWARE TOOL

The DDP process involves gathering and reasoning with quantitative data on requirements, failure modes, PACTs, and relationships between these. Software tool support for this process is essential. The first step in this direction used spreadsheets as the implementation, successfully demonstrating proof of concept of the process, and establishing the directions in which more customized support was needed. An implementation effort to fulfil these needs was initiated two years ago, and versions of this implementation have been used in the DDP applications described in this paper. This section describes the key aspects of this software tool support, its realization in the current implementation, and our continuing development efforts.

### *DDP software tool support needs*

To be effective in complex domains such as spacecraft technology, the DDP process must combine multiple experts' knowledge. To date, applications of the DDP process have been organized as sessions in which those experts contribute their knowledge, explore the implications of its combination, and make decisions on selection of PACTs, etc. The needs of the software tool that supports these sessions are therefore to:

- Accommodate on-the-fly input of DDP knowledge
- Perform the numerical calculations that underpin the combination of the quantitative knowledge
- Render the information in suitable visualizations
- Facilitate the experts in exploring alternatives (selections of PACTs) and record their decisions

The DDP software tool was developed to satisfy these needs, and we continue to evolve it as our understanding of these needs improves. The primary challenge faced by this software is to gather a non-trivial amount of information and make it available in ways that support human understanding. The DDP tool is *not* used simply to gather data followed by an automatic determination of the optimal solution. In part this is because there are many factors that need to be taken into account, and an efficient way to conduct a DDP application is to have the experts perform the decision making, aided by the DDP tool.

### *Realization in implementation*

The DDP implementation takes the form of a Visual Basic program, using an Access database for permanent storage. It runs as a stand-along program on Windows platforms. Where appropriate, elements of its look-and-feel replicate features common to the Windows style of interface, augmented by DDP-specific aspects.

The tool offers a small number of key "views", suited to the portrayal, entry and editing of the various forms of DDP information. These are:

- Tree views – used for *hierarchies* of requirements, failure modes and PACTs. Like the Windows File Explorer, large hierarchies can be handled by expanding and collapsing subtrees. Individual elements, and subtrees of these elements, can be turned "on" and "off", via checkboxes (again, making use of a familiar interface mechanism). The primary purpose of this capability is to allow users to turn on and off PACTs, in their search for a judicious selection from among them. It also facilitates "what if" studies (e.g., "what if we were to give up this requirement?"), and allows rapid customization of generic knowledge (e.g., turn off a failure mode that has no relevance to the case at hand). During DDP sessions, users may (and do) enter new elements into these trees, search through them, and restructure them by copy & paste, drag-and-drop, and promotion and demotion of subtrees.

- Matrix views – used for the quantitative relationships between requirements and failure modes, and between PACTs and failure modes. Like simple spreadsheets, header rows and columns are used to display the titles of the elements (e.g., requirements), while the inner cells display the numerical values of the quantitative relationships (e.g., degree of impact of a failure mode on a requirements).

**Figure 8** The requirements view in the DDP tool.... Note that the view for Failure Modes is similar (bar char is now of risk levels), the view for PACTs is also similar (bar chart is now of PACT Figure of Merit)

- Bar chart views – the DDP implementation automatically computes the combined impact of failure modes on requirements, and effectiveness of PACTs on failure modes. The results of these computations are displayed in bar charts.

    - The "Requirements Drivers" bar chart displays the status of each of the requirements – how much they are impacted by failure modes. Each bar is subdivided into the degree to which its requirements was originally at risk (prior to the selection of any PACTs), and the degree to which it is currently at risk taking into account the mitigating effects of the currently selected set of PACTs. Choosing the option to have these in sorted order reveals the requirements most at risk.
    - The "Risk Balance" bar chart displays the status of each of the failure modes – how much impact they are causing to the requirements. Each bar is subdivided into the degree to which it originally impacted requirements (prior to the selection of any PACTs), and the degree to which it is currently impacting requirements, taking into account the mitigating effects of selected PACTs. Choosing the option to have these in sorted order reveals the failure modes contributing the most risk.
    - The "PACTs" bar chart displays the status of each of the PACTs – how much impact-to-requirements savings its selection would achieve, subdivided into the savings it would accomplish independent of the other PACTs (i.e., as if the other PACTs were all unselected), and the savings it would accomplish above and beyond the impact savings already achieved by the currently selected PACTs. Choosing the option to have these in sorted order reveals the PACTs that have the most risk reducing effects.

- Compact list views – used as an alternative to matrices for the display of the quantitative relationships. Modelled after "stem-and-leaf charts" (a style which Tufte attributes to John W. Tukey, "Some Graphical and Semigraphic Displays"), they are a compact way to portray DDP's relatively sparse matrices.

These various views are kept automatically synchronized, so that as the user makes an update via one such view, the others are automatically updated accordingly. Over the course of time we have added features that facilitate navigating the complex risk landscape through these various views.

The tool generates printable reports, and saves the inputs to, and results of, a DDP session into a database, which can be re-opened at a later date for further scrutiny, extension, etc.

While primarily a stand-alone implementation, we have built some import/export capabilities to interact with companion tools that also manipulate risk-related information, such as the AskPete cost estimation and planning tool from the NASA Glenn facility[8].

*Ongoing implementation efforts*

We are engaged in an effort to pre-populate DDP with knowledge bases of failure mode and PACT information, gathered from experts in the various domains pertinent to spacecraft development. We expect this will necessitate extending the tool, to better facilitate downselecting from the relatively large body of information to the subset pertinent to the task at hand.

We are also looking into adding capabilities to automatically searching for near-optimal sets of PACTs. This obviously relies on the resource costs of PACTs to have already been provided. In DDP applications to date, this information has not been elicited for entry into the tool, and therefore PACT selection has been a primarily manual process. In this current practice, the tool presents cogent visualizations of, essentially, just the benefit side of the cost-benefit balance of PACT selection.

The Visual Basic implementation is the one we have used in support of the DDP process applications so far. We are continuing to maintain and extend this implementation for the type of DDP sessions described throughout this paper. We[††] are also pursuing the development of a Java-based implementation, which is being architected to accommodate some extensions and variations on the DDP process, especially:

- More elaborate logical structures between and among DDP elements, e.g., failure modes organized into fault trees, whose nodes are "and" or "or" nodes with the usual logical connotation; PACTs whose application would have the side effect of inducing additional failure modes.
- User-customizable computations (e.g., users may choose to adjust the formula used to calculate the combined impact of multiple Failure Modes that all impact the same requirement).
- Collaborative and/or distributed DDP sessions, in which multiple users may each have their own view into the DDP space of information (through their own computer terminal), and may each contribute to the information. Effectively coordinating the on-the-fly merging of inputs from multiple participants is expected to be a critical issue to this mode of operation.
- Increased interoperability, through compatibility with data interchange standards (e.g., XML, STEP).

[††] The Java development is being done by Dr. Julia Dunphy.

## APPLICATION TO TECHNOLOGY PRODUCTS

### *Technology Infusion Challenges*

It has been widely recognized that a technology "gap" exists between the NASA R&D community and the focused flight project and mission communities. Interaction between the two communities is still not done well and this has exacerbated an already problematic NASA technology infusion problem‡‡. The goal of this application of the DDP process is to increase the infusion rates of advanced technology into NASA flight missions by uniting the two communities and retiring risk associated with technology infusion.

We expect many technology developments to "die on the vine" at lower maturity levels where product viability is (by the nature of early R&D) unknown. This is the nature of aggressive research programs and "dead ends" can still be looked at as successes. However, historically we have assumed that once technology viability has been demonstrated, the remaining "engineering" work needed to make the product useful, should be easily achieved in a flight project development environment. Unfortunately, what we see instead is that overlooked failure modes, undetected or not considered earlier in the product validation phase, surface at the worst possible time when ongoing projects are counting on the technology to be ready in time for their mission. Furthermore, while these failure modes are predominately "simple" engineering issues these tend to be show stoppers because they are discovered so late in the project development life cycle. Technology infusion rates suggested by a recent survey which looked at a random sample of technology "pull" situations, indicate that once a product has reached proof-of-concept (POC), it stands less than a 40 percent chance of infusion into a flight system. With the cost and visibility of flight projects being so high, this has created a skeptical mission community which will usually seek less risky COTS solutions at the first sign of trouble with a given new NASA technology. The R&D community which is unfamiliar with traditional flight project practices, gain little education in the process of technology infusion into flight systems and there is no real mechanism to manage the technology gap.

### *DDP Impact on technology infusion*

This has led us to develop a new way of handling risky technologies which utilizes the DDP process in a series of small workshops, and involves engaging a small team of multi-disciplined "experts", from a wide variety of fields to "trouble shoot" the technology well before a project engages to use it in space applications. Also participating in the workshop are the researchers developing the product and the specific project/mission customers who provide and negotiate requirements. Risk issues and risk mitigation techniques are weighted by the team and an approach is agreed on to address each potential failure mode over the course of the infusion process. A final "risk balancing" step selects (and deselects) work thus forming the final risk portfolio and technology roadmap.

Technology developers who have participated in say they gained a clear understanding of how their technology "fits" into the defined application, and found strengthened customer advocacy since there was now a clear workable plan. In addition, the contacts made in the workshop could now help them in the transition into flight. A third benefit is that they gained a strong position in the proposal competition since a multidisciplinary panel containing their customers all agreed to the technology infusion progression. . Often it was seen that substantial efforts in non-technology development areas were necessary to sufficiently mature the product prior to technology pull, and funding could now be sought (and justified) in non-R&D programs to get this work done. The opposite was true as well where it was realized that the R&D efforts currently being worked were not necessarily the highest priority, and a course correction was necessary.

To date, we have applied DDP to four component level developments and one software development with extremely successful outcomes. In one case, a the clarification of a major customer requirement led to ~$1.2M savings in work not required, and a product delivery to the customer two years earlier. In another case, a technology targeted for termination due to a lack of customer interest, and poor hope for success, was rescued and is now proceeding to multiple customer utilization. The software study led to consideration of a commercial software development environment to replace the expensive software design practices used at NASA today. Another technology development was discovered to be a hopeless waste of funding given it's progress, status, and team attrition situation.

## CONCLUSIONS AND FUTURE WORK

The DDP process has been described. The process (and associated software realization) is intended to enable life cycle risk management. The process asks users to think about where they want to be (requirements), what could get in the way (failure modes) and what can be done about the obstacles (PACTs). The DDP process then asks for scoring of impacts and effectiveness and produces a number of outputs which are intended to aid in the achievement of continuous risk management. The DDP process is a systematic, top-down approach which integrates bottoms-up information. Through requirement, failure mode and PACT tree evolution, the process utilizes all available

‡‡A series of informal conversations with a variety of industrial partners would appear to indicate that the industrial community at large could benefit from an improved technology infusion process.

information to provide the most up-to-date view of the risk landscape. By combining both intuitive and analytical information, the DDP process can be implemented over the entire project life cycle.

While a large number of applications of the DDP process have already been performed, it is the goal of this year's work to begin implementation on an entire project for it's entire life cycle. Several candidates have expressed interest but other opportunities are welcome.

The technology evaluation efforts to date have led to a JPL initiative to institutionalize the process and evaluate all JPL technologies which have reached 'proof of concept' and to assist in developing a roadmap for likely flight infusion. The next step in application to the technology infusion process is to develop the infrastructure to take the information collected, and develop a "skunkworks" technology maturation process to accelerate risk retirement in any given technology, outside of the project and mission environments. This process will be focused on making robust products out of proven concepts, prior to project PDR. In this way, the transition from R&D to CDR will finally be managed completely.

We are willing to team with other organizations to develop additional add-ons for the tool and encourage beta-testing of the tool to generate additional feedback regarding utility and possible future directions.

The tool development work continues and this year is focused on developing an automatic optimization engine, implementing across the web, interfacing to other existing tools and adding additional user features. In addition, this year we are developing wizards which will attempt to couple to institutional expert opinion, closed loop process and corrective action evaluations, lessons learned, and other knowledge bases.

*Acknowledgements*

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

*Biographies*

***Steven Cornford*** *graduated from UC Berkeley with undergraduate degrees in Mathematics and Physics and received his doctorate in Physics from Texas A&M University in 1992. Since coming to JPL he focused his early efforts at JPL on establishing a quantitative basis for environmental test program selection and implementation. As Payload Reliability Assurance Program Element Manager, this evolved into establishing a quantitative basis for evaluating the effectiveness of overall reliability and test programs as well as performing residual risk assessments of new technologies. This has resulted in the Defect Detection and Prevention (DDP) process which facilitates the risk evaluations which are the subject of this paper. He received the NASA Exceptional Service Medal in 1997 for his efforts to date. He has been an instrument system engineer, a test-bed Cognizant Engineer and is currently involved with improving JPL's technology infusion processes. He is also currently the Principal Investigator for the development and implementation of the DDP software tool, including assisting JPL/NASA technologists in implementing a process for evaluation of early technology development efforts.*

***Martin Feather*** *obtained his BA and MA degrees in mathematics and computer science from Cambridge University, England, and his PhD degree in artificial intelligence from the University of Edinburgh, Scotland. For a number of years Dr. Feather worked on NSF and DARPA funded research while at the University of Southern California's Information Sciences Institute. He is currently a member of the Technical Staff at NASA's Jet Propulsion Laboratory. His research interests encompass the early phases of software development, especially requirements engineering and analysis/V&V. He is the programmer for the tool shown in this paper.*

***Kenneth Hicks*** *holds a B.S. in Computer Science and Engineering and joined NASA's Jet Propulsion Laboratory in March of 1985 in order to pursue a career developing advanced avionics for space and other critical terrestrial applications. He began his career developing advanced VLSI sensors for space environment and spacecraft health monitoring. Since then, he has served as the Cognizant Engineer for the Clementine Radiation Reliability and Assurance Experiment (RRELAX) which was flown successfully in 1994. Ken has also managed a task to develop an Electronics Test Bed (ETB) for the BMDO-sponsored STRV-2 spacecraft, which comprises five separate technology demonstrations and environment monitors. Each of these tasks led to Exceptional Achievement awards. He also served as Cognizant Engineer for the X-33 Avionics Flight Experiment to be launched on X-33 in 2000. Ken is now the Implementation Technologist for the NASA Cross-Enterprise Technology Development Program, where he is now spending significant time investigating various infusion mechanisms for eventual use across the NASA Centers, and is attempting to improve NASA's advanced technology infusion rate.*

## REFERENCES


[1] T. Gindorf and S. Cornford, "Defect Detection and Prevention (DDP): A Tool for Failure Mode Risk Management", *presented at the 50th IAC, Amsterdam, The Netherlands, October, 1999.*

[2] S. Cornford, "Managing Risk as a Resource using the Defect Detection and Prevention process", *International Conference on Probabilistic Safety Assessment and Management,* September 13-18, 1998.

[3] M.S. Feather, S.L. Cornford, and M. Gibbel, "Scalable Mechanisms for Requirements Interaction Management", *in Proceedings, 4th IEEE International Conference on Requirements Engineering, Schaumburg, Illinois, 19-23 June 2000. IEEE Computer Society*

[4] For licensing information, contact Jennifer Schlickbernd at JPL at (818)354-2241. The tool is available via a no fee licence to NASA Centers and their subcontractors working on NASA tasks.

[5] Steven L. Cornford Mark Gibbel Martin Feather David Oberhettinger , "A Physics/Engineering of Failure Based Analysis and Tool for Quantifying Residual Risks in Hardware", *Proc. Ann. Reliability & Maintainability Symp., 2000 Jan.*

[6] see NASA NPG 7210.5, "Program and Project Risk Management Processes and Requirements"

[7] Assurance Effectiveness Guidebook, edited by Mr. Tim Larson and prepared by the NASA Failure Detection and Prevention Program

[8] http://tkurtz.grc.nasa.gov/pete/